DOCUMENT RESUME

ED 328 485 SO 021 069

AUTHOR Baker, Michael
TITLE Tutoring with Incomplete and Uncertain Knowledge. CITE Report No. 38.
INSTITUTION Open Univ., Walton, Bletchley, Bucks (England). Inst. of Educational Technology.
PUB DATE May 88
NOTE 40p.; Paper presented at the European Seminar on Intelligent Tutoring Systems (Tubingen, West Germany, September 1987).
PUB TYPE Reports - Descriptive (141) -- Speeches/Conference Papers (150)

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Artificial Intelligence; *Cognitive Processes; Computer Assisted Instruction; Computer Uses in Education; Foreign Countries; Knowledge Level; Learning Processes; *Musical Composition; Music Education; *Programed Tutoring; Teaching Methods
IDENTIFIERS England

ABSTRACT

The design of an Intelligent Tutoring System (ITS) in a knowledge domain where expertise is modeled as a set of uncertain and incomplete beliefs that are justifiable and expressible in the form of a critical argument is outlined. Issues concerning knowledge communication in a tutorial interaction are discussed with reference to a cognitive model for perception of musical structures, used in an ITS for structural analysis and interpretation of tonal music. A tutorial dialog planning system is outlined that uses sets of context spaces (Reichman 1986) whose discourse function matches possible tutorial actions. Nine figures and a 62-item bibliography are included. (Author/DB)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

SCOPE OF INTEREST NOTICE

The ERIC Facility has assigned this document for processing to: SO

In our judgment, this document is also of interest to the Clearinghouses noted to the right. Indexing should reflect their special points of view. IR

ED 268 645 SO 021 045

CENTRE FOR INFORMATION TECHNOLOGY IN EDUCATION

IET

TUTORING WITH INCOMPLETE AND UNCERTAIN KNOWLEDGE

MICHAEL BAKER

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

X This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
C. J. LAWLESS
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

CITE Report No. 38

SO 021 045

Institute of Educational Technology
The Open University
Walton Hall
MK7 6AA
Tel. 0908 74066

TUTORING WITH INCOMPLETE AND UNCERTAIN KNOWLEDGE

Michael Baker

May 1988

This paper was presented at the European Seminar on Intelligent Tutoring Systems, in Tübingen (West Germany), during September 1987. It will appear in a special issue of the International Journal of Educational Research, on 'Knowledge Acquisition in Intelligent Learning Systems', during 1988.

Michael Baker
Institute of Educational Technology
The Open University
Walton Hall
Milton Keynes MK7 6AA
telephone (0908) 653488
email MJ_BAKER@UK.AC.OPEN.ACS.VAXA

May 1988.

Abstract

An architecture is described for an Intelligent Tutoring System in a knowledge domain where expertise is modelled as a set of uncertain and incomplete beliefs, which are justifiable and expressible in the form of a critical argument. Issues concerning knowledge communication in a tutorial interaction are discussed with reference to a cognitive model for perception of musical structures, used in an ITS for structural analysis and interpretation of tonal music. Justifications can be generated from the system's beliefs by representing the concepts which interrelate possible justification types, but we conclude that such justifications must be set within a tutorial dialogue which aims to satisfy the multiple goals of teacher and learner. A tutorial dialogue planning system is outlined, which uses sets of *context spaces* (Reichman 1986) whose discourse function matches possible tutorial actions.

Introduction 1

Modelling and Communicating Knowlege of Musical Structures

Modelling music cognition 2

Communicating beliefs concerning musical structures 3

Teaching concepts without confirming or disconfirming instances 4

Knowledge Communication and Tutorial Dialogues

Explanation and dialogue 7

The structure of critical argument 9

Planning tutorial dialogues - a first pass 11

High level tutorial goals 11

Dialogue coherence 13

Representing and reasoning about beliefs 14

The dialogue planner 15

Summary and Conclusions 16

Acknowledgements 18

Bibliography 19

Figures

Figure 1 - The system's knowledge of musical structure 25

Figure 2 - Justifying possible instances of 'tea'. 25

Figure 3 - The space of possible justifications for instances of 'tea'. 26

Figure 4 - Possible justifications for phrase boundaries. 27

Figure 5 - Toulmin's structure for reasoned argument. 28

Figure 6 - The dialogue planning architecture. 29

Figure 7 - A high level tutorial goal. 30

Figure 8 - A lower level 'Socratic tutoring' goal. 30

Figure 9 - Hierarchy of tutorial goals. 31

Introduction

Existing research in Intelligent Tutoring Systems ('ITS') has mostly concentrated on highly formalised domains such as physics, mathematics and computer programming, where the nature of the knowledge domain can often support a highly directive form of tutoring which concentrates on imparting factual knowledge. This paper describes ongoing ITS research in the unusual domain of analysing the structure of tonal music. Consideration of such domains may well radically revise our attitude to knowledge representation and tutorial strategies in an ITS because of the necessity to deal with knowledge which is uncertain, incomplete and admitting several possible acceptable solutions. Deciding on the most appropriate way of modelling human knowledge in a domain is possibly the most important factor in ITS design. We would argue that the design of an ITS for the analysis of musical structure, where such models have yet to be developed, provides a good opportunity to consider afresh how knowledge of the domain should be modelled in a way which facilitates its communication in a tutorial interaction.

The paper briefly describes a program which analyses musical structures, based on a theory of human music cognition. We discuss the way in which choice of modelling technique was influenced by the necessity to derive this knowledge in a form which is cognitively valid and hence communicable in a tutorial interaction. Starting from the set of uncertain beliefs and their justifications derived by the model, we describe how 'explanations' could be generated by including an additional knowledge-base in terms of a semantic network (Quillan 1968) which interrelates justifications and the types of assertions which they can justify. After reviewing the literature on explanation, we conclude that generation of explanations as individual utterances is not sufficient basis for a tutorial interaction, and that explanation needs to be seen as an emergent property of an ongoing cooperative tutorial dialogue. A "discourse genre" approach to generating critical arguments based on use of schemas is considered unsuitable for use in tutorial contexts. A brief description is given of future research directions in developing an ITS architecture which can support such tutorial dialogues. It is based on planning (and replanning) sets of possible utterances of the system and learner, viewed in terms of their dialogue focus and function.

Modelling and Communicating Knowledge of Musical Structures

Modelling Music Cognition

The level of uncertainty involved in musical grouping perception appears to be such that it is only possible to attempt to delimit or constrain the space of possible grouping structures, and to specify weak decision mechanisms for preferring one or more of a number of competing hypotheses. In the system which we describe here ('INTERPRET'), a program called GRAF ('GRouping Analysis with Frames') has been developed to simulate human grouping perception, and has been described in detail elsewhere (Baker 1988a, Baker, forthcoming 1989). Our present discussion emphasises the way in which the modelling approach chosen was intended to facilitate knowledge communication in a tutorial interaction.

Grouping analysis in music presents similar problems to those faced by speech understanding systems, such as HEARSAY-II (Erman et al 1980), and generally has the style of integrating a number of contributing knowledge sources on the 'musical surface' (large leaps , rests and other changes) and at a deeper level of unconscious inference, in terms of harmonic and other structures. In designing this part of the ITS a number of possible approaches were considered and rejected in terms of the depth and cognitive validity of the knowlege representation derived. For example, approaches which depended on the use of production systems with certainty factors or weightings attached to rules were rejected here as embodying a level of knowledge representation which is not deep enough to form a model of what Wenger (1986) terms "communicable knowledge". The approach adopted combines a psychologically plausible parser with a system of schemata (implemented as frames, Minsky 1977) which encapsulate stereotypical grouping structures used for expectancy-based parsing. The grouping frames are represented in a hierarchical and heterarchically connected system of *context spaces*, (Grosz's term, 1979) which encapsulate the listener's knowledge of specific musical genres, and the overall real-time parsing process uses techniques which have been proposed for modelling the processes which doctors use in diagnosing diseases (Pauker, Gorry, Kassirer & Schwartz 1984). For our purposes here, what is important is that the system has to rely on some kind of *plausible reasoning*, where it may settle too early for solutions which may turn out to be incorrect, or may fail to find any solution at all. Usually, the system derives two or three solutions concerning the structure of the melody, in terms of frames which record the positions of structural boundaries, together with justifications for these beliefs in terms of factors used during the course of the computation. For example, one frame hypothesis might record a belief that there was a phrase

boundary after the first eight notes of the melody, which was justified in terms of an interval leap and recurring rhythmic units on either side of the boundary. We would maintain that this uncertainty in the knowledge-base of the system is not a deficiency, but that it adequately represents the true nature of human understanding, which may be unclear, incomplete, uncertain and even completely absent: a psychologically plausible parser *must* fail sometimes.

In the next section we consider how this knowledge-base of beliefs concerning musical structure can be represented in an educationally communicable way.

Communicating Beliefs concerning Musical Structure

Intelligent tutoring systems usually embody several representations for the domain which they aim to communicate. At the level of knowledge representation described so far we have a set of *justified beliefs* concerning possible phrase-boundaries present in the melody. In addition the system has a set of principles for preferring one belief+justifications set over another, and a set of global constraints on sets of phrase boundaries which could be instantiated at any one time in a melody. Diagrammatically, it could be represented as in **Figure 1**. For example, we might have

b_1 [belief] = 'phrase-boundary at point p1 in the melody'

$j_1 1$ [justification] = 'exact match in harmony [vi IV ii] in note-sets n1 &n2'

$j_1 2$ = 'exact match, 3 bars, reductional level 3'

b_2 = 'phrase-boundary at point p2 in the melody'

$j_2 1$ = 'exact match in harmony [vi IV ii] in note-sets n1 &n2'

$j_2 2$ = 'exact match, 3 bars, reductional level 2'

A preference principle (prefer a reductional match at the highest level')applies here, to prefer b_2, since level 2 is higher than level 3. In addition, supposing that a boundary close to rb_2 had already been asserted, a global constraint (prefer boundary sets which create structural units of equal size), could be invoked.

A crucial aspect of learning to justify possible instances of concepts appears to be an understanding of the *concepts* and beliefs upon which justifications depend, since no acquired concept exists in a 'conceptual vacuum'. The highest level representation in the system is thus intended to facilitate acquisition of these concepts which underlie beliefs concerning musical structure. Since the knowledge of instances of the concepts which our system aims to communicate is limited and uncertain, this precludes, for example, the majority of theories of concept learning, which depend on providing

definite examples and 'near misses' of concepts (Winston 1977) . We shall describe an approach which aims to facilitate learning of concepts in the domain, and which seems to follow naturally from its representation. A simplified tutorial situation concerning structural analysis of music might be as follows: A student may believe that a structural division termed a 'phrase boundary' occurs after the first n notes of the melody being analysed, and may justify this in terms of similarities in harmony on either side of the boundary, and the fact that there is a rest of two beats at this point. The tutor [system] may have a justified belief that the phrase boundary in this section of the piece is elsewhere, and will attempt to facilitate the acquisition of the concept of what a phrase boundary is by engaging the student in a critical discussion which is appropriate to the student's current understanding of the domain.

Teaching Concepts Without Confirming or Disconfirming Instances

In order to motivate the proposed approach we shall consider an everyday example, which may be more intelligible to most readers than more detailed musical examples. Suppose you want to teach the concept of what 'tea' is, to a someone who already has some intuitions concerning what tea should be like, but you lack a ready stock of instances of cups of tea. You do, however, have a cup of possibly tea-like liquid in front of you, which you have reason to believe is tea, and your friend has a similar cup of liquid. We can view learning the concept 'tea' under these circumstances, as resulting from a dialogue in which your friend advances justifications for believing that they have identified an instance of that concept, and you have incomplete knowledge of instances of tea, but know the space of possible justifications which are relevant to something being an instance of tea or not. When presented with available justifications, your friend can state reasons for believing that this is an instance of tea, which can be represented diagrammatically as in **Figure 2.** You may want to explain other kinds of evidence which are relevant to the question of whether this is tea or not, but in order to do this, you want to be sure that your friend will be able to understand those additional justifications. Now assume that your knowledge can be encapsulated in an ITS. The system would represent the concepts which underlie alternative justifications, from its own point of view. From a previous conversation, the system has a model of which concepts the student does and does not understand. We can represent the system's knowledge diagrammatically in **Figure 3**, including the student's stated justifications and conceptual knowledge (via student model, shaded).

The system can generate justifications for instantiating the concept 'TEA' in the absence of any definite instances, by critiquing the student's justifications with respect to those known to the system. The idea is that once the student is aware of further factors which are relevant to the concept's instantiation, (s)he will be able to grasp the concept. Rather than simply stating a long list of the possible justifications, the tutoring system will take account of the concepts which underlie those justifications, in terms of conceptual knowledge attributed to the student. The idea is a simple and intuitive one: *the system should attempt to relate its explanation of justifications for the concept to the user's existing knowledge and interests.* It can use the semantic network (termed a 'conceptual layer') to do this. For example, the teacher might relate the explanation to the user, initially by beginning with the justifications which they have mentioned, and working via available conceptual links to related justifications. If a conceptual link is 'missing' (i.e. the system has no knowledge that the student is familiar with that concept) then this generates a possible tutorial goal for the teacher. Taking the example above, since the student has stated that the colour of the liquid is relevant to whether it is tea, the tutor can start from this statement and decide to explain the justification "hot" (i.e. that tea is a liquid which is usually drunk when hot). Connecting links are concepts of "physical_property" and "temperature", and since these have not been explained before, and the teacher does not believe that the student is aware of these, these two concepts may need to be tutored. A human teacher might give the following minimal explanation: "The physical properties of these kinds of liquids are usually relevant to telling what they are. You mentioned the colour; but the fact that it's also hot is another relevant physical property, since tea is a liquid which is usually drunk when hot."
It is not difficult to see how, using this kind of representation, a teacher could present other justifying evidence which was relevant to instantiation of a concept, using further conceptual links, nor is it difficult to see how a computer-based tutor could generate something like this explanation using simple formats or templates.
Figure 4 shows part of the conceptual layer attached to the concept 'phrase_boundary', known to the system, which could be used to generate justifications in a similar manner to that described for the 'toy' example.

In considering this example, we assume that at some point in the teaching interaction, the system decides that this is a good time to tutor the concept phrase. This decision would depend on an overall tutorial planner, which constructs tutorial subgoals to teach certain concepts into a coherent plan, a user-model which has knowledge of the concepts believed to be known to the user and their connectivity, and interaction knowledge of the current focus

of attention, and so on. Justifications could thus be generated in a technically quite simple manner. Taking the example where both the student and system believe that 'phrase_boundary' is instantiated at the same point p, the justification could be generated as follows:
[1] fetch justification-pairs which match exactly between the system and student justification-sets. For these alternatives, the system has of course already determined their factual accuracy, so look in the conceptual layer which attaches to each justification, and check that the student understands underlying concepts up to the highest node in the layer, with reference to the student-model. For 'missing' concepts, generate a new tutorial subgoal of tutoring that concept for the tutorial planner.
[2] fetch the set of justifications put forward by the student, not recorded by the system, and check their factual accuracy [note that the editor implicitly constrains the user from advancing justifications which are not of the kind which are possible for this concept, in the system's terms]. For each justification, find the system's justification[s] which are closest to this - share a common node at lowest level above - and explain how these alternatives known to the system are also possible in virtue of the common conceptual links. If these links are not known to the student, generate a new tutorial subgoal, as for [1]. Continue until there are no paired system/student similar justifications remaining, then do [3] with the single remaining student or system justification for which there is no such close link.
[3] check the conceptual layer above the justification, and generate tutorial subgoals as appropriate. If it is a student justification, then look for other places in the music where this <attribute><value> pair is also true, and present this to the student (they might want to revise their opinion that this is still a justification here). If it is a system justification, present it only if the student possesses sufficient concepts to understand it, according to the student model (assuming an available reliable modelling technique). Otherwise, generate tutorial goals of tutoring these concepts.

We have only given a summary sketch of how justifications could be constructed using the conceptual knowledge underlying alternative justifications. However, it is clear that the generation of such separate justifications is not sufficient to define the overall structure of a tutorial interaction. In the next section we review recent work in explanation and conclude that an explanation can not simply be identified with generation of a set of utterances.

Knowledge Communication and Tutorial Dialogues

Explanation and dialogue

Any ITS whose set of beliefs concerning a knowledge domain is uncertain and possibly incomplete clearly should not attempt to simply impart that knowledge to the student, since those beliefs are debateable and possibly incorrect. A reasonable alternative is to develop a system which is able to articulate the reasoning processes which led to its acquisition of those beliefs. One approach to understanding how an ITS might help students to learn to justify their own beliefs is to view the problem as one of the system *explaining* its reasoning and beliefs.

A large body of research in explanation has focussed on the problem of how to add explanation facilities to Expert Systems. There are a number of ways, however, in which most attempts to add explanation facilities to so-called 'second generation' expert systems are inadequate as a basis for adaptation to tutor novices in their domain of expertise. Common approaches have been to find some way of 'pruning' the information already present in the system's chain of rule-based reasoning (Weiner 1980), or to generalise rules to embody more abstract concepts (Clancey 1987).The approach of 'pruning' chains of rules may assist in preventing stating what is obvious or already said to *an expert*, but such explanations are of little use in tutoring the domain to non-experts, who will not understand the concepts embodied in those rules. Such generation of explanations gives the user no opportunity to interact and *negotiate* the form of the explanation required beyond an isolated utterance, and it does not explain the reasoning processes by which conclusions are derived. In a sense, this is fortuitous, since the reasoning methods used by expert systems are certainly not those which human experts would use and hence should not be inculcated into students.

An approach to tailoring an interaction with an expert system to a user's interests is to *critique* the user's approach with respect to the system. An example is Miller's (1984) ATTENDING system, which critiqued physicians' plans for courses of anaesthetic management. A critique is a kind of 'reasoned explanation' of the significant differences between two belief sets, or in this case, treatment plans. From the point of view of intelligent tutoring, critiquing systems as they stand do not help novices to construct plans in the first place, nor do they aim to teach the domain concepts to be embodied in such plans. It seems clear that in order to do this, a more extended *tutorial dialogue* would have to take place, in terms of which we can view a critique as a form of dialogue which is constrained to a single

utterance from the system and the user. In a tutorial context, a single response to a student's solution would rarely be sufficient: the student would usually want to respond to the tutor's critical response, perhaps to question a factual assumption, request clarification on a conceptual matter, and eventually to submit a revised solution on the basis of negotiation.

The 'Explanation and Dialogue' sections of the past three Alvey Workshops on Explanation (see references by Draper and by O'Malley) have reached the conclusion that *explanation requires an extended dialogue*, rather than generation of isolated utterances. The principal reasons why a more fully interactive dialogue is required concern: (i) establishing the knowledge assumed by participants in the explanatory dialogue, (ii) establishing 'relevance constraints' (what will be accepted as an explanation by all participants?), and (iii) the necessity for establishing mutual goals. In accordance with our earlier description of semantic layers and their use in generating justifications, Draper (1987b) argues that human explanatory dialogues depend on a 'differencing' operation being carried out between the belief-sets of explainer and inquirer:

> "In terms of a [semantic] network model of mind, a request for an explanation is a request for E [the explainer] to show I [the inquirer] part of E's network".

Other authors have argued for a view of explanations as dynamically following their own course, guided only by a set of local dialogue focus rules and global coherence constraints (O'Malley 1987). From the point of view of constructing explanatory dialogues in an ITS, however, the views expressed by Draper and O'Malley are reconcilable. If one accepts that, to a greater or lesser extent, *tutoring dialogues have a goal of facilitating knowledge communication*, then we must assume Draper's position concerning the differencing mechanisms in explanation. Whilst tutoring dialogues may have a set of long term teaching objectives, this is not incompatible with their pursuit via the interaction of local discourse rules, which do not assume any underlying fixed schema (O'Malley 1987). O'Malley also argues that explanation does not only (if at all, in her terms) involve negotiation to establish mutually held beliefs, but involves negotiation of alternative and possibly conflicting viewpoints on the subject of the dialogue (an issue which has been adressed in ITS research by the attempt to model and utilitse alternative viewpoints on knowledge domains (Wenger 1987, p.355 ff)).

We conclude that in order for our proposed system to tutor with uncertain knowledge it must engage the student in an extended dialogue in the form of a *critical argument*. In the next section an approach to constructing tutorial dialogues in the form of a critical argument is described. Its (somewhat ambitious) educational purpose is broadly twofold:

(1) *critical argument*
The system should engage the student in a dialogue which models the structure of human reasoning.The objective is to teach the activity of forming critical arguments by engaging the student in that activity.
(2) *concepts in knowledge sources*
Forming a critical arguent in a domain involves not just knowledge of how arguments may be developed, but also of what knowledge sources are required to support such arguments, in terms of a set of underlying concepts.

The resulting tutorial dialogue should combine the structure of critical arguments with a dialogue structure based on tutoring domain concepts involved in justifications advanced in support of arguments. There is a further constraint on the kind of dialogue which the system must generate: in a tutorial context we do not simply want a coherent dialogue, but also one which achieves *tutorial goals*. We shall argue that in order to do this, the system would need to incorporate a dialogue planner.

The Structure of Critical Argument

One possible approach to computational generation and understanding of sustained argument is to specify some normative or schematic dialogue structure. Significantly, a body of this research has been carried out within the knowledge domain of archaeological reasoning, a humanity which shares many important properties with arguments over structural analysis of music.

The most recent and complete exposition of this approach is Stutt's (1987) 'Arguing Expert System' (the Argument Support Program for Archaeology , ASPA), which builds on earlier work on argument structure by Alvarado, Dyer and Flowers (1986). The ASPA system (Stutt 1987) uses a limited number of argument patterns or scripts in an attempt to extend the explanation capabilities and user acceptance of second generation expert systems. As an interaction of argument exchanges proceeds, the system builds a dependency graph, based on Toulmin, Rieke & Janik's (1979) analysis of sound human reasoning. **Figure 5** shows a segment of one side of an argument exchange, and the dependency graph would consist of a network of these and other structures, based on parsing user input into elements of these script-like structures.

The motivation behind using such argument scripts in expert system research is to give a greater symmetry between the operations on the knowledge base which are available to the user and system, thus improving

user-acceptance. The argument exchanges may thus proceed in a way which lacks directedness, which is perfectly acceptable for the purposes for which the approach is intended. In a tutorial context, however, we need to take account of the goals of both the tutor and the learner, and integrate these goals within the overall dialogue. The kinds of activities which occur in critical arguments in a tutorial context include many features, such as informing, advising, correcting, and so on. We would therefore argue that for tutorial purposes, it would not be fruitful to consider the dialogue in terms of inflexible structures, and that it may not be useful to label the exchange as a 'critical argument'. The dialogue would certainly include elements of what we mean by argument, but should be able to flexibly accomodate the way in which a dialogue may alternate between criticism, explanation, statement, question, Socratic questioning and any other form of interaction which is viewed as achieving mutual goals of tutor and learner. The approach of identifying schematic argument structures is an example of an approach to discourse analysis which Reichman (1986) terms that of identifying "discourse forms" (p.9). Her criticism of the "genre dictates structure" approach is as follows:

> "Since it is clear that in any ordinary discourse or conversation, speakers use a number of different "genres" - stories, descriptions, arguments, for instance - what conventions govern the development of the whole conversation that includes these subparts? An adequate theory of discourse would need to specify the conventions or rules that apply across as well as within discourse forms..."
> (Reichman 1986, p.9).

Stutt's arguing expert system considers the dialogue in terms of a microstructure of paired argument exchanges. We would argue that tutorial dialogues must do justice to larger scale dialogue structures, including dialogue focus (Grosz 1979, Reichman 1986) and larger scale goals of conversants (Power 1979, Apppelt 1983).

An ITS which aims to engage the learner in critical argument which is educationally beneficial, should have powerful reasoning methods concerning the user's knowledge and likely way of participating in the dialogue, the dialogue should be coherent and focussed, and should attempt to achieve the mutual goals of system and learner. In the next section we describe a system based on planning an interaction based on the system's and user's possible sets of utterances, considered as essentially goal-directed *speech acts* (Searle 1969), to achieve high level tutorial goals.

Planning Tutorial Dialogues - a first pass

Planning the course of a dialogue depends on exploiting those aspects of linguistic utterances which can be viewed as species of *actions*, and are termed *speech acts*. We shall not attempt to describe or review the extensive literature on Speech Act Theory, but the following are some of the most important areas of research. Searle (1969) has synthesised and extended Speech Act Theory, building on earlier work by Wittgenstein (1953) and Austin (1962). Authors such as Power (1979) and Appelt (1983) have used techniques derived from planning in AI robot systems to study discourse structure and the generation of goal-directed multi-agent dialogues. See Suchman (1985) for an influential critique of the use of plans in human-machine communication.

In this final section we now describe an approach to planning tutorial dialogues. Our objective in specifying this architecture is to find a way of structuring the system's set of beliefs and those which it ascribes to the student into a coherent dialogue which achieves multiple goals.**Figure 6** below shows an overview of the proposed system.

We assume that we have some reasonably powerful method for reasoning about the student's beliefs (user-modelling), and that the tutorial interaction is at a stage where we have ascribed a number of beliefs to the student from an up-to-date dialogue history. We shall briefly describe the function of each section of the proposed dialogue planning architecture.

High level tutorial goals

At the highest level, the tutorial goals are to teach concepts of musical structure, and such concepts are ordered in a hierarchy of goals and subgoals. It should be mentioned, however, that there are other goals which are not represented or taught explicity, which are not directly concerned with tutoring concepts, but rather with teaching the student to use such concepts as part of a coherent argument to express their own point of view concerning musical analysis. Intuitively, the idea is that we want to teach the student to argue their case as part of a coherent dialogue, where the means of teaching this is to attempt to engage them in the activity of arguing itself. We do not propose to address the very difficult problem of attempting to infer the function in dialogue of natural language utterances made by the student, but rather to use a constrained interface with which the student must make an assertion, justification or warrant, indicating explicitly (menu selection) the believed role of the utterance, and its connection to the

system's previous utterances. For example, a student may state an alleged justification which is not of the kind which can justify a particular kind of assertion (in terms of its set of possible grounds). An important tutorial objective would be to help the student to understand the way in which their utterances could function in argument. For purposes of exposition, we shall omit these goals from a first pass at specifying tutorial goals.

As a first prototype we propose to use a simple STRIPS-like representation for goals, preconditions, actions and their assumed effects (Fikes & Nilsson 1971), whilst retaining the possibility that future prototypes may use an hierarchical planner (Sacerdoti 1977). In terms of tutorial dialogues, we aim to plan to achieve changes in models of student beliefs, using a sequence of tutorial actions to achieve student and teacher goals. Tutorial actions include interactions via a graphic music editing tool, as well as generation of sets of utterances. As an initial simplification, we assume complete co-operativity in the dialogue in terms of joint goals and plans. Our sequence of STRIPS operators would take the form shown in **Figure 7**. At lower levels in the goal hierarchy, goals would be similar to the work of Collins and Stevens on interactive teaching (Collins & Stevens 1983) human plausible reasoning (Collins 1978, Collins & Michalski 1987) and the structure of Socratic dialogues (Stevens & Collins 1977). **Figure 8** shows examples of lower level goals which aim to generate the kinds of counter-examples which a Socratic tutor might put forward. Note that for each operator, each 'slot' need not take a value - a goal may have no preconditions or actions, and may simply expand down to a series of subgoals which may or may not themselves possess all elements. The diagram in **Figure 9** shows the form which such a hierarchy of goals might take, where associated actions would call appropriate local discourse procedures which operated on the semantic net associated with those claims. At this stage we shall not attempt to specify the precise search strategy to be used in interpreting a plan - usually the precise combination of means-ends analysis and goal regression required depends on specific properties of the problem addressed. Fundamentally, these high level tutorial goals constitute the top-down input to the planner, which would result in a plan for a sequence of illocutionary acts to achieve specific goals. For example, we might have a sequence of illocutionary actions to achieve the goal 'teach phrase boundary' in the form:

INFORM student of other places where the stated justification-type applies, but where they have not postulated a phrase-boundary

REQUEST response from student in terms of possible alteration of their beliefs

AGREE with student response

Attached to these illocutionary acts, we propose to use local discourse procedures (see Elsom-Cook 1984, 1985, 1987b, Power 1979) to generate sets of surface speech acts, as opposed to planning single utterances. The bottom-up input to the planner takes the form of the possible assertions available to both the tutor and the student (the student model) at a particular point in the dialogue. We briefly discuss the form which these utterances take in the next section.

Dialogue coherence

Below the level of illocutionary acts, the top-down element of the planner must use the set of utterances available to the system in terms of the knowledge-base derived by the music parsing system. Grouping these utterances solely in terms of actions to achieve a plan could, however, result in a dialogue which appeared unconnected, which lacked *local coherence* although it may possess some global coherence in virtue of participating in a set of planned illocutionary acts . We therefore need a way of characterising the topic or focus of sets of possible assertions, so that we can use a set of rules for dialogue focus shifts and coherence to act as a constraint on the planner's selection of a set of utterances to achieve the current goal. An abstract record of previous foci in the dialogue will be maintained as a *dialogue history*. Since we aim to plan in terms of the student's possible utterances in terms of their beliefs, we need to either make that assumption that they will have a similar conception of dialogue coherence, or that their conception can be modelled with a set of rules which are variants of those used by the system. Given our statement that we want to teach them how to argue critically and coherently, we could view it as an educational virtue of the system if it constrained the student to use a conception of coherence which the system aimed to impart.

We propose to exploit Reichman's (1986) work on discourse focus and context to specify *coherence constraints* on the planner's use of a set of assertions in performative utterances. At the heart of Reichman's work is the notion of a *context space*, being an abstract discourse constituent which is essentially delimited by a specific set of utterances performing a conversational move. *Conversational moves* are formalised as a discourse grammar for coherent dialogue, between context spaces. It is in terms of the *function* of a set of possible assertions and justifications defined as a context space (eg, SUPPORTING , CHALLENGING, INTERRUPTING) that the system's knowledge can be organised and accessed by the planner. The planner will match top-down tutorial goals into a form equivalent to function of context spaces in the system's and user's beliefs. For example, at one planning level we might have a goal, 'INFORM <student s1> <counterexample>, which

would connect with a set of system beliefs , defined as a context space whose function matches the goal. These context spaces will be called on by discourse procedures acting on the concept representation to generate specific utterances. A grammar of 'conversational moves' can then be used to ensure coherence between the planned sequence of focus shifts performed by the user's and system's utterances, given some means of determining focus (formal algorithms for determining dialogue focus and tracking its movement have been provided by Sidner, 1986). We shall discuss what might happen if the system is not able to elicit the kind of utterance which it needs from the student in order to execute its tutorial plan in the next section.

The kinds of conversational moves which Reichman describes are intuitively acceptable, and many related directly to dialogues which involve elements of criticism and argument. For example,

> "... we cannot have purported statements of support without the presence of a claim to be supported, nor can we have an explanation without something to be explained."
> (Reichman 1986, p.24).

An example is the SUPPORT conversational move, which includes a set of standard subparts, such as the claim, its principle of support (warrant), relevant supporting facts, and so on. The SUPPORT move is usually followed by a RESTATEMENT of the claim, which acts as a conversational clue that the SUPPORT move is over. Use of a conversational move can close the preceding context space, or put it in controlling status, as well as having other possible effects. An additional important set of utterances which need to be planned are those which communicate these shifts of focus.

Representing and reasoning about beliefs

We discussed an additional knowledge base required by the system to integrate possible justifications and assertions in terms of a particular kind of semantic net. We termed this a 'concept layer', since its role is to characterise the assumed concepts which link possible justifications to the kinds of assertions which they may correctly justify. In terms of reasoning about the student's beliefs we propose no new user-modelling technique (since that is not the major focus of this work) but propose to explore the use of Appelt's (1983) notion of representing an agent's mental states as possible worlds of belief sets, and using Moore's (1980) integration of this approach into first-order predicate calculus to reason about these possible mental states.

The dialogue planner

Given a list of STRIPS operators, a dialogue model, student and system beliefs and a set of dialogue coherence rules, the plan interpreter which would integrate these inputs into a dialogue plan involves a number of complex problems in planning, some of which are as yet unachieved research goals. Clearly, we shall not attempt to propose a solution to these problems here, but simply discuss how they arise in the area of planning tutorial dialogues.

Most work in conventional planners involves constructing a plan for a single goal, which can generally be solved by the application of a large quantity of world knowledge in terms of actions to achieve goals and their presumed effects on the world. However, in the real world context of a flexible dialogue, a 'third person planner' (i.e. not one of the participants in the dialogue) has to deal with satisfation of multiple goals simultaneously. It is not difficult to think of situations in which the pursuit of one goal may prevent the satisfaction of another conjunctive goal since goals may interact in numerous complex ways. A number of solutions have been proposed for dealing with interacting plans, in the work of Sacerdoti (1977), Sussman (1975), Tate (1975) and Warren (1974), all of which contain explicit mechanisms. One approach is to develop a plan to achieve one goal, then pass the other conjunctive goals over that plan to attempt to modify it to achieve the others as well (Waldinger 1977). In Sacerdoti's (1977) approach, a hierarchical plan is developed top-down from a highly abstract goal, with a collection of 'critics' examining the plan at each level of expansion to guard against possible conflicts. At this stage we simply record the fact that there are a number of possible available techniques which address these problems.

Most planners make use of what has become known as the 'STRIPS assumption', which is simply to assume that actions leave the world unchanged other than for those effects which we explicitly represent. Clearly, this will not always be the case within the space of possible coherent dialogues, so we need mechanisms for determining the effects of actions (diagnostically testing the effects of tutorial actions on the state of the student model) and replanning where necessary. We may also need to replan in cases where we cannot elicit an appropriate coherent response from the student and where we want to deal with interruptions which are not relevant at the present but would be at some later possible point in the plan. Checking the result of plan actions is often termed 'execution monitoring' , where "There is as yet no general theory of execution monitoring (Charniak & McDermott, 1985 p.524). Again, there are a number of existing techniques

for replanning without completely abandoning an existing plan, and for doing localised 'fixups' to preserve an existing plan. Clearly, it would require quite complex decision procedures to decide which of these alternatives should be pursued.

Summary and Conclusions

Our stated initial objective was to indicate the research directions which could be followed to develop an ITS which could deal with uncertain and incomplete knowledge domains. We have made a first pass at indicating an approach within the domain of analysing musical structures based on planning dialogues which embody critical argument within general goal-oriented dialogue.

Clearly, the fact that the system has no undisputable facts at its disposal means that it can not simply use subset student modelling and highly directive tutorial strategies in attempt to make the learner into a copy of the tutor. We argue that a suitable response is to engage the student in a critical argument, to both teach the structure of how to coherently express their position, draw attention to any inconsistencies, and make students aware of the kinds of knowledge sources which are relevant to critical arguments which could be adduced in favour of their position. Structural analysis of music is a knowledge domain which forces us to address the problem of critically arguing for a point of view, but it is not confined to this domain. Clearly, much of what practising scientists and researchers do is to argue their position with colleagues in the form of a human dialogue, rather that in a set of mathematical or logical proofs. Dialogues involving critical argument are obviously central to humanities such as politics, philosophy and debating, law and archaeology (Stutt 1987).

We argued that mere generation of explanations or justifications from a knowledge base was not sufficient for tutorial purposes, and that explanation should be viewed as an emergent property from a sustained dialogue. The kinds of tutorial dialogues which we have described can not be simply categorised into the 'genre' of arguments, explanations or justifications: we argued that critical argument in a tutorial context must be viewed as dialogue *per se*. One way of trying to generate such dialogues is to view sets of utterances defined by their dialogue function as possible goals for which a planning system might generate a set of actions. The planning system should attempt to achieve the multiple goals of the ITS's tutorial goals, with the perceived learning goals of the student, constrained by a set of

dialogue coherence rules between discourse foci. We stated a number of problems which such a planning system would face, in terms of interactions between multiple goals, planning in a changing world, and in the possibly unpredictable effects of tutorial actions on the student's mental states.

We have not proposed solutions to all of these problems. Nevertheless, we have clearly indicated an approach which an ITS could adopt in knowledge domains which lack clear and complete facts, and where the ability to express beliefs in the form of a coherent critical discourse is an important skill.

Acknowledgements

My warmest thanks to my supervisor, Dr. Mark Elsom-Cook, who has generously contributed many of the ideas on tutorial dialogue in this paper, and given me the will to continue. Dr. Claire O'Malley has given me her time and ideas on explanation, together with much needed encouragement. Thanks to Arthur Stutt for explaining his research to me and letting me have the benefit of his knowledge about the structure of critical argument. Fiona Spensley kindly read the final draft, made a number of useful criticisms, and provided many essential cups of coffee. I am grateful to Professor Tim O'Shea for providing initial encouragement, opportunity and stimulus for this work. I owe more than I can say to Susan's love and care during my labours.

Advice on typographic design was kindly provided by Rachel Hewson, but all responsibility for the final appearance of this document is my own.

Bibliography

Alvarado, S.J., Dyer, M.G. & Flowers, M. (1986). Editorial comprehension in OpED through argument units. Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86), Philadelphia, PA, 250-256, 1986.

Appelt, D.E. (1982). Planning Natural-Language Utterances to Satisfy Multiple Goals, Technical Note 259, March 1982, SRI International, California 94025.

Austin, J.L. (1962). How to Do Things with Words, Oxford University Press 1962.

Baker, M., Burstein, M.H. & Collins, A.M. (1987). Implementing a Model of Human Plausible Reasoning. Proceedings of IJCAI 1987, p. 185-188.

Baker, M.J. (1988, in press). A Cognitive Model for Perception of Musical Grouping Structures. forthcoming in Contemporary Music Review, Winter 1988, special issue on *Music and the Cognitive Sciences*.

Baker, M.J. (1988, in press). An Artificial Intelligence Approach to Musical Grouping Analysis. forthcoming in Contemporary Music Review, Spring 1989, special issue on *Science and Music*.

Carbonell, (1970). AI in CAI: an artificial intelligence approach to computer-assisted instruction. IEEE Transactions on Man-Machine Systems, vol.11, no.4, pp.190-202.

Charniak and McDermott (1985). Introduction to Artificial Intelligence. Addison-Wesley 1985.

Clancey,W. (1987). Knowledge-Based Tutoring: The GUIDON Program, MIT Press 1987.

Cohen, P.R. (1985). Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach, Pitman 1985.

Collins, A., Warnock, E.H., & Passafiume, J.J. (1974) Analysis and Synthesis of Tutorial Dialogues, BBN Technical Report no. 2789, March 1974.

Collins, A., Warnock, E.H., Aiello,N. & Miller, M.L. (1975). Reasoning from

Incomplete Knowledge. In Representation and Understanding, Collins, A. & Bobrow, D.G. (eds.), Academic Press 1975.

Collins, A. (1978). Fragments of a Theory of Human Plausible Reasoning. In Proceedings of the Conference on Theoretical Issues in Natural Language Processing-2 (TINLAP-2), 194-201, New York: ACM, 1978.

Collins, A. & Michalski,R. (1987, in press). The Logic of Plausible Reasoning: A Core Theory. Cognitive Science, in press.

Collins, A. & Stevens, A.L. (1983). A Cognitive Theory of Inquiry Teaching, in Instructional-Design Theories and Models: An Overview of their Current Status, (ed.) Reigeluth, C.M., Lawrence Erlbaum, 1983.

Davey, A. (1978) Discourse Production. Edinburgh University Press 1978.

Draper, S. (1987a). Report on the Dialogue Section of the 2nd Alvey Workshop on Explanation, Proceedings of the 2nd Alvey Workshop on Explanation, University of Surrey (Great Britain), January 8-9th, 1987.

Draper, S. (1987b). A User-Centred Concept of Explanation, Proceedings of the 2nd Alvey Workshop on Explanation, University of Surrey (Great Britain), January 8-9th, 1987.

Draper, S. (1987c). What's going on in everyday explanations. To appear in C. Antaki (ed.) Analysing lay explanations: a casebook of methods, (Sage Publications, London, in press).

Draper, S. (1987d). Spontaneous Explanation: report on the dialogue section of the 3rd Alvey workshop on explanation. Proceedings of the 3rd Alvey Workshop on Explanation, University of Surrey (Great Britain), 29th September-1st October, 1987.

Elsom-Cook,M. (1984). Design Considerations of an Intelligent Tutoring System for Programming Languages, 1984. PhD thesis, Warwick University (Great Britain).

Elsom-Cook, M. (1985). Towards a framework for human-computer discourse. In Humans and computers: Designing the Interface, eds. Johnson & Cook, Cambridge University Press, 1985.

Elsom-Cook (1987) Discourse for Tutoring. Paper presented at the 3rd International Conference on Artificial Intelligence and Tutoring,

University of Pittsburgh LRDC, May 8-10th 1987.

Erman, L.D., Hayes-Roth, F. Lesser, V.R. & Reddy, D.R. (1980). The HEARSAY-II Speech-Understanding System: integrating knowledge to Resolve Uncertainty. ACM Surveys vol 12 no.2, June 1980.

Fikes & Nilsson (1971) STRIPS: A new approach to the application of theorem proving to problem solving. Proceedings of the Second International Joint Conference on Artificial Intelligence, Imperial College London, 1st-3rd September 1971.

Flowers, M., McGuire, R. & Birnbaum, L. (1982). Adversary arguments and the logic of personal attacks. In Strategies for Natural Language Understanding, (eds.) Lehnert, W.G. & Ringle, M.G. , Lawrence Erlbaum Ass., NJ, 1982.

Grice, H.P. (1975). Logic and conversation. In Syntax and Semantics Volume 3: Speech Acts, (ed.) Cole,P. and Morgan,J.L., Academic Press 1975.

Grosz, B.J. (1979). The representation and use of fucus in dialogue understanding. PhD Dissertation 1979.

Johnson-Laird, P.N. & Wason,P.C. (1977) Thinking: Readings in Cognitive Science, Cambridge University Press 1977.

Langlotz, Curtis P., & Shortcliffe, Edward H., (1983) Adapting a consultation system to critique user plans, Int.J. Man-Machine Studies, (1983) **19**, 479-496

McGuigan, S. & Black, J.B. (1986). Creation and Comprehension of Arguments. In Knowledge Structures, Galambos, J.A. & Black, J.B. (eds.) Lawrence Erlbaum, 1986.

Medlin, D.L. & Smith, E.E. (1984). Concepts and concept formation. Annual Review of Psychology, 35: 113-38.

Miller, Perry L. (1984) A Critiquing Approach to Expert Computer Advice: ATTENDING, Pitman Publishing 1984.

Minsky, M. (1977). Frame-system theory. In Johnson-Laird, P.N. & Wason, P.C. (eds.) Thinking: Readings in Cognitive Science, Cambridge University Press 1977.

O'Malley, C. (1987). Applying Studies of Natural Dialogues to Human-Computer Interaction: Are Two Heads Better than One? Proceedings of the 2nd Alvey Workshop on Explanation, University of Surrey (Great Britain), January 8-9th, 1987.

Moore, R.C. (1980). Reasoning about knowledge and action, SRI International AI Center Technical Report No. 191, 1980.

Pauker, S.G., Gorry,A., Kassirer, J.P. & Schwartz, W.B. (1984). Towards the Simulation of Clinical Cognition: Taking a Present Illness by Computer. In Readings in Medical Artificial Intelligence, (eds.) Clancey, W.J. & Shortcliffe, E.H., Addison-Wesley 1984.

Petrie-Brown, A. (1987) The Influence of Context and Coherence as a Foundation for Dialogue Research, Open University CITE Report no. 8 , February 1987.

Power, R. (1979) .The organisation of purposeful dialogues. Linguistics 17 (1979), 107-152.

Quillan, M.R. (1968) . Semantic Memory. In Minsky, M. (ed.) Semantic Information Processing, Cambridge Mass., 1968, pp. 227-270.

Reichman, R. (1985) Getting Computers to Talk Like You and Me, MIT Press, Cambridge (Mass.), 1985.

Sacerdoti , E. (1977). A Structure for Plans and Behavior. Elsevier North-Holland, Amsterdam.

Schank, R.C. (1986). Explanation Patterns: Understanding Mechanically and Creatively. Lawrence Erlbaum 1986.

Searle, J. (1969). Speech Acts, Cambridge University Press, 1969.

Self, J. (1986). The Application of Machine Learning to Student Modelling. Instructional Science **14** (1986) 327-338.

Sidner, C.L. (1986). Focussing in the Comprehension of Definite Anaphora. In Readings in Natural Language Processing, (eds.) Grosz, B.J., Sparck-Jones, K. & Webber, B.L., Morgan Kaufmann 1986.

Steels, Luc (1987). The Deepening of Expert Systems. Artificial Intelligence Communications, vol.0, no.1, August 1987.

Stevens, A.L. & Collins, A. (1977). The Goal Structure of a Socratic Tutor, BBN Technical Report no.3518, March 1977.

Stevens, A.L., Collins, A. & Goldin, S. (1978). Diagnosing Students' Misconceptions in Causal Models, BBN technical report no. 3786, June 1978.

Stutt, A. (1987) Second Generation Expert Systems, Explanations, Arguments and Archaeology, Human Cognition Research Laboratory Technical Report No. 25, December 1987, The Open University, Milton Keynes, Great Britain.

Suchman, L.A. (1985). PLANS AND SITUATED ACTIONS: The Problem of human-machine communication. XEROX PARC technical report P85-00005, Xerox Corporation, Palo Alto Research Centers, 3333 Coyote Hill Road, Palo Alto, California 94304.

Sussman, G.J. (1975). A Computer Model of Skill Acquisition. American Elsevier, New York.

Tate, A. (1975). Interacting goals and their use. Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, Georgia, USSR.

Toulmin, S., Rieke, R. & Janik, A. (1979) An Introduction to Reasoning. Macmillan Publishing Co., New York, 1979.

Waldinger, R. (1977). Achieving Several Goals Simultaneously. Machine Intelligence 8, (eds.) Elcock, E.W. & Michie, D., Ellis Horwood Ltd., 1977.

Warren, D.H.D. (1974). WARPLAN: A system for generating plans. Department of Computational Logic, Memo no. 76, University of Edinburgh.

Weiner, J.L. (1980). BLAH, a system which explains its reasoning. Artificial Intelligence **15** (1980), 19-48.

Wenger, E. (1987) .Artificial Intelligence and Tutoring Systems: Computational and Cognitive Approaches to the Communication of Knowledge, Morgan Kaufmann Inc. , 1987.

Wilensky, R. (1983) Planning and Understanding. Addison-Wesley, Reading

Mass., 1983.

Winston, P. (1977) Learning to identify toy block structures, in Thinking - Readings in Cognitive Science, eds. Johnson-Laird, P.N. & Wason, P.C., Cambridge University Press 1977.

Wittgenstein, L. (1953). Philosophical Investigations. Translated by G.E.M. Anscombe, Basil Blackwell, Oxford 1953.

Worden, R.P., Foote, M.H., Knight, J.A. and Andersen, S.K. (1986). Co-operative Expert Systems. Proceedings of the 7th European Conference on Artificial Intelligence (ECAI), July 1986, pp. 319-334.

Figures

Figure 1

Summary of the system's knowledge of musical structure, derived from the cognitive model of musical grouping perception.

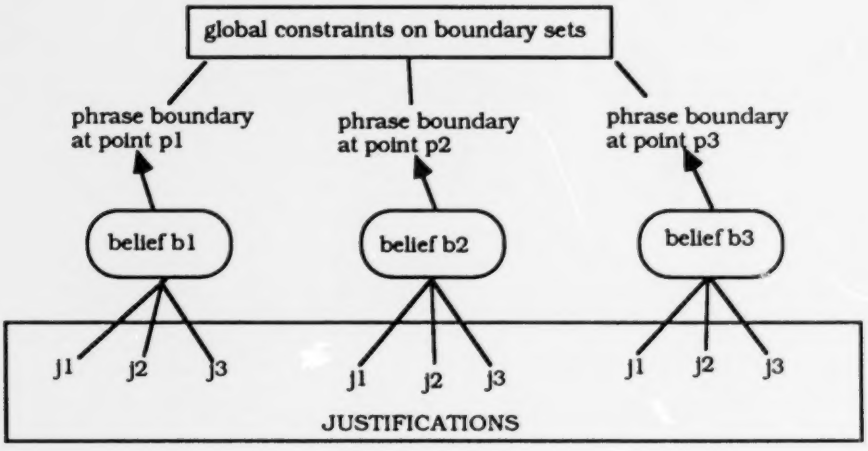

Figure 2

Possible justifications for believing that a liquid is tea.

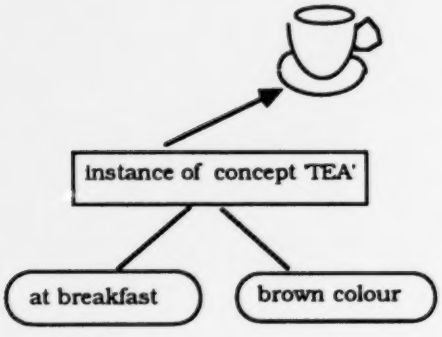

Figure 3

Representation of the space of possible justifications for instantiating the concept 'TEA', with associated 'conceptual layer'. Concepts which the system believes the student to possess are shaded.

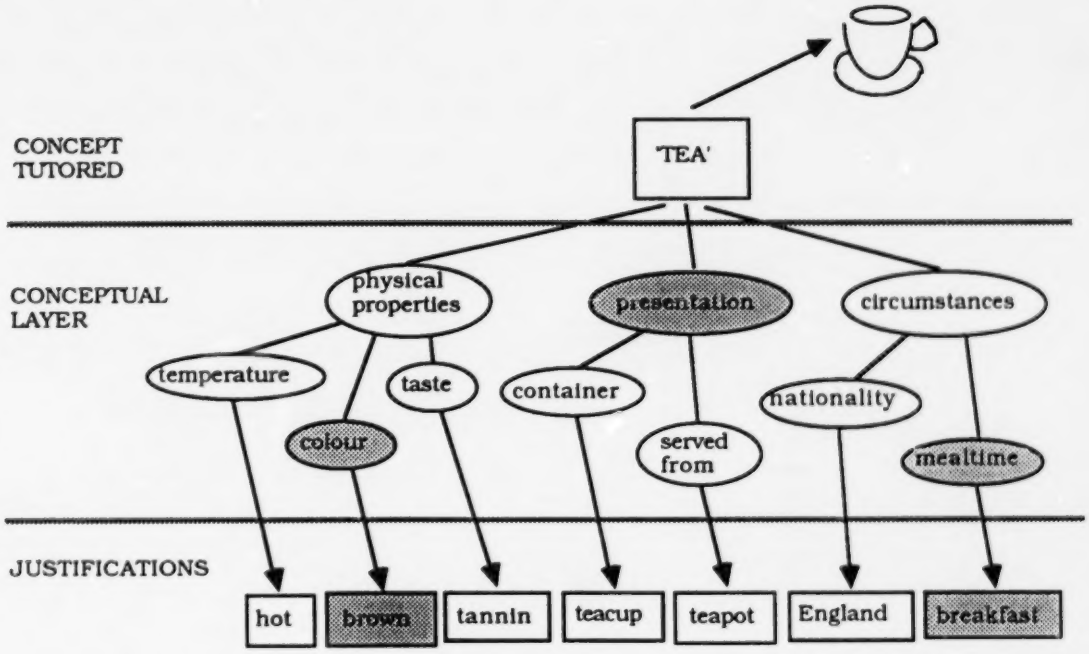

Figure 4

Some possible justifications for instantiating the concept 'PHRASE BOUNDARY'.

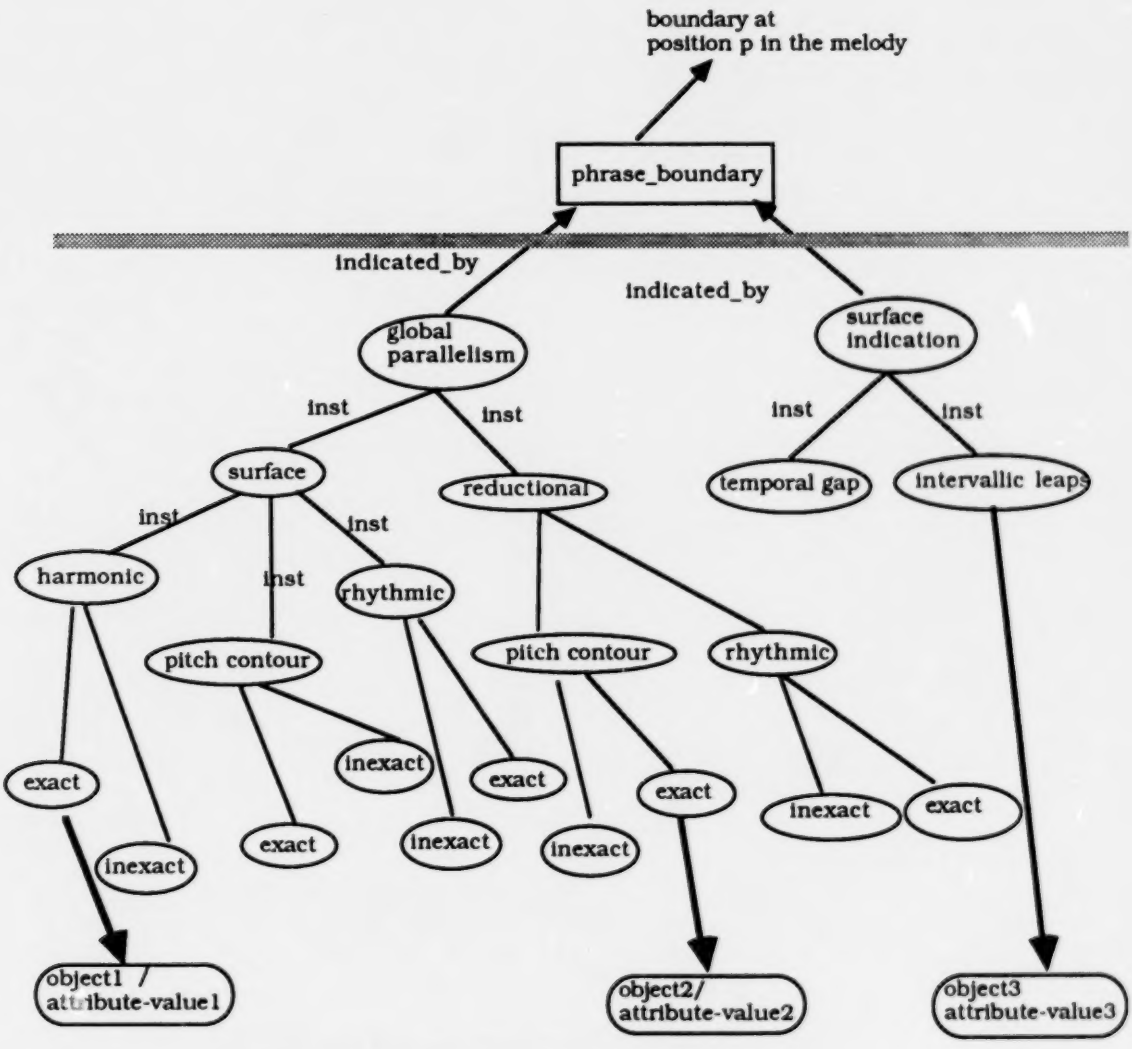

EXAMPLE: a phrase boundary is suggested at point p in the melody by
[1] an exact match in a possible harmonisation for 3 bars on either side ;
[2] an exact match in pitch contour for 4 bars on either side at reductional level3;
[3] an interval leap of an octave, when notes move in step on either side.

Figure 5

Toulmin et al's (1979) structure for human reasoned argument.

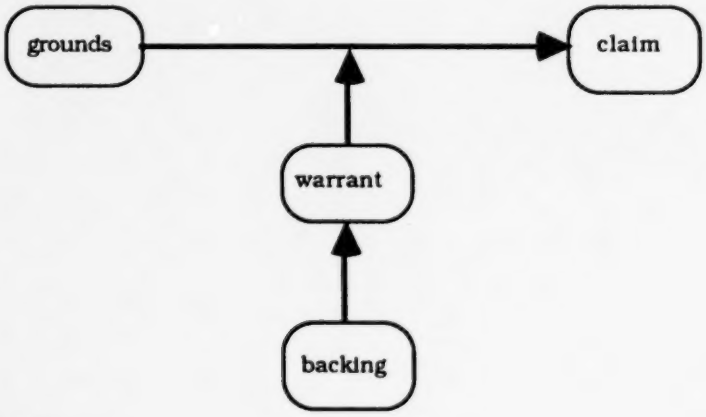

EXAMPLE

CLAIM: your computer is on fire!
GROUNDS: there's smoke coming out of the disk drive
WARRANT: there's no smoke without fire
BACKING: (science has shown?) that smoke can generally only be produced by combustion

Figure 6
Overview of the dialogue planning architecture.

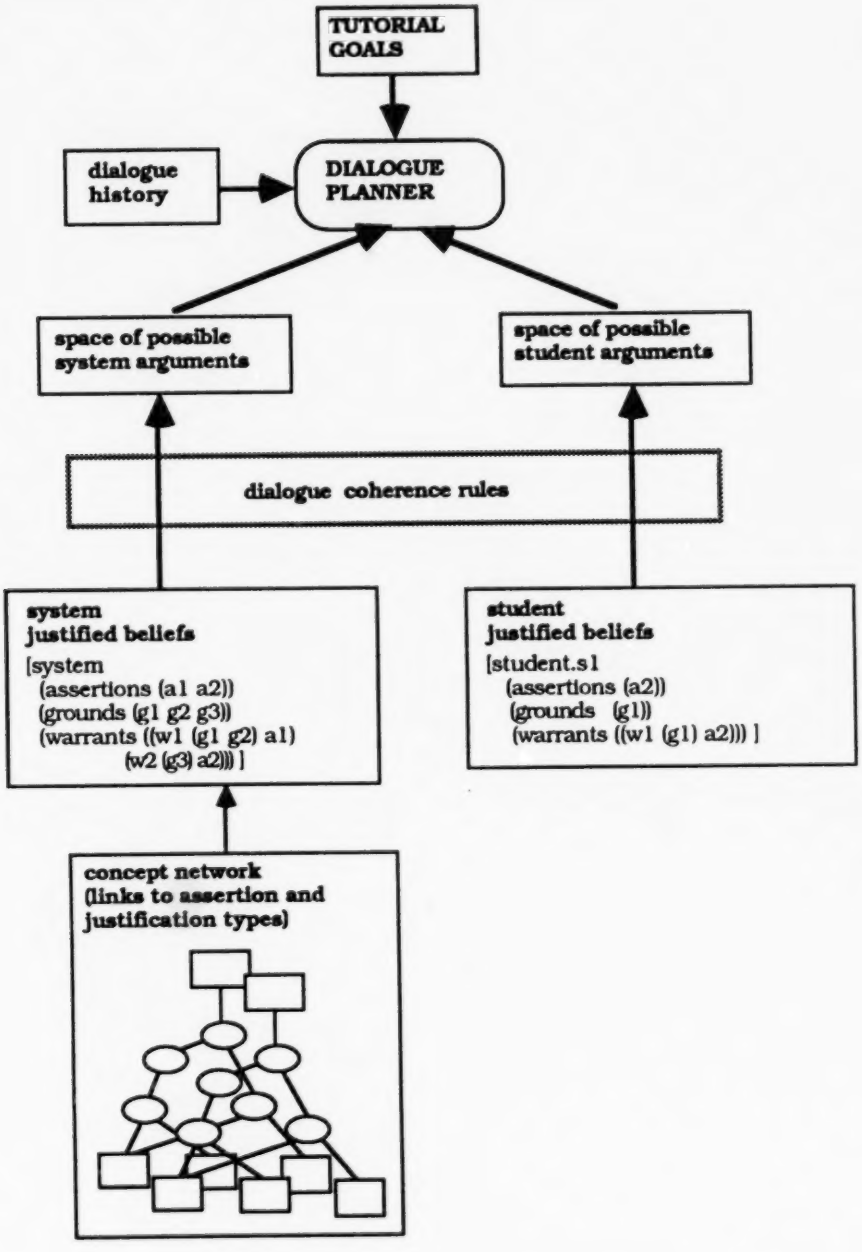

Figure 7
An example of a high level tutorial goal, represented in a STRIPS-like form.

goal	teach musical structure
preconditions	not knows (student) musical structure
subgoals	teach concepts (c1 c2 c3 c4) & review concepts (c1 c2 c3 c4).
effects	knows (student) musical structure
actions	nil

Figure 8
Example of a lower level tutorial goal, relating to goals of a Socratic tutor (Collins & Stevens 1983).

goal	show counter-example
preconditions	common phrase-boundary belief student belief subset of tutor's exists non-phrase-boundary with student belief
subgoals	identify non-phrase-boundary elicit student prediction comment student prediction
effects	none
actions	none

Figure 9

A possible hierarchy of tutorial goals.

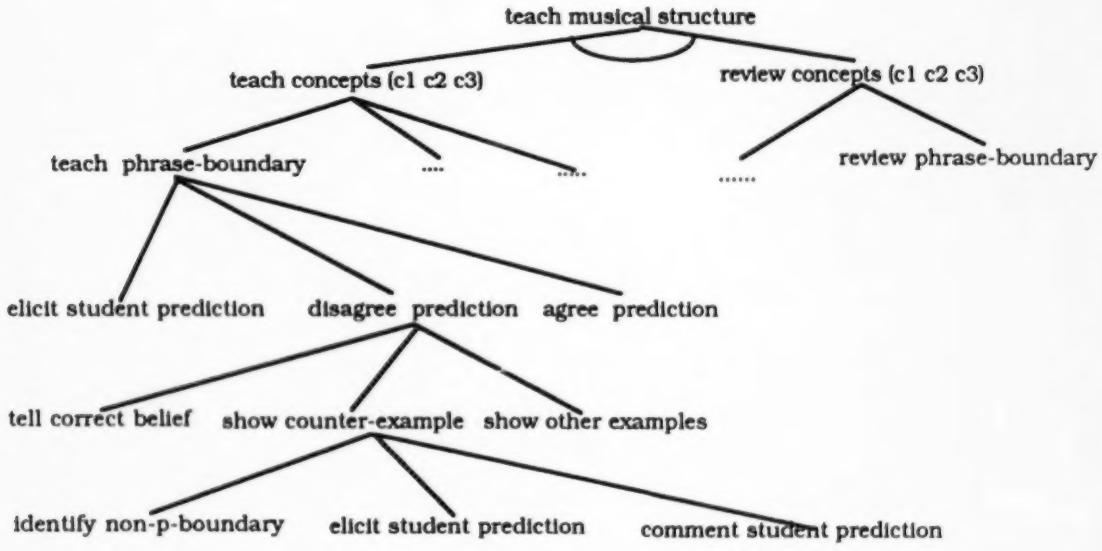

CENTRE FOR INFORMATION TECHNLOGY IN EDUCATION

List of CITE Reports

These reports may be obtained from:

Hansa Solanki, Institute of Educational Technology, The Open University, Walton Hall, MILTON KEYNES, MK7 6AA, England.

Report No.	Title and Author
1	A.T. Vincent, (1985) Computing and the Blind.
2	A. Jones, G. Kirkup, J. Morrison (1985) A Trial of Home Based Computer Terminals.
3	Gill Kirkup, (1985) The Present and Potential Use of Ceefax in the Open University.
4	Mark Elsom-Cook, (1986) Artificial Intelligence and Computer Assisted Instruction.
5	Mark Elsom-Cook, (1986) A Pascal program checker.
6	Simon Holland, (1986) How computers are used in the teaching of music and speculations about how Artificial Intelligence could be applied to radically improve the learning of composition skills.
7	Simon Holland, (1986) Design consideration for a human-computer interface using 12-tone three-dimensional harmony space to aid novices to learn aspects of harmony and composition.
8	Alison Petrie-Brown, (1987) The Influence of Context and Coherence as a Foundation for Dialogue Research.
9	Eileen Scanlon, Randall B. Smith (1987) A Rational Reconstruction of a Bubble Chamber Simulation Using The Alternate Reality Kit.
10	Mark Elsom-Cook, (1987) Intelligent Computer-Aided Instruction research at the Open University.
11	Mark Elsom-Cook, (1987) Towards a framework for human-computer discourse.
12	Mark Elsom-Cook, (1987) MATILDA AND IMPART: Lisp tools.
13	Mark Elsom-Cook, (1987) Guided discovery tutoring and bounded user modelling in Intelligent Computer Aided Instruction ed. J. Self, Chapman-Hall 1987.
14	A.M. Petrie-Brown and M.T. Elsom-Cook, (1987) An Examination of an AI model of indirect speech acts.
15	A. Edwards, (1987) Integrating Synthetic Speech With Other Auditory Cues In Graphical Computer Programs For Blind Users.
16	S. Holland, (June 1987) A knowledge-based tutor for music composition.

17 S. Holland, (June 1987) New Cognitive Theories of Harmony Applied To Direct Manipulation Tools for Novices.

18 M. Baker. (July 1987) Intelligent Computer-Aided Instruction and Musical Performance Skills.

19 M. Baker, (August 1987) Proposed Research Directions for Intelligent Computer-aided Instruction in Musical Performance Skills.

20 A D N Edwards, (August 1987), Adapting interfaces for visually disabled users.

21 M Elsom-Cook, (September 1987), Acquisition of computing skills.

22 M. Baker, (September 1987), Computational Analysis of musical grouping structures.

23 M Baker, (June 1987), Automated Analysis of Musical Grouping Structures as a Basis for a Guided Discovery Environment for Interpretation of Music.

24 D Laurillard, (October 1987), The different forms of learning in psychology and education.

25 A.D.N. Edwards, (November 1987) Modelling blind users" interactions with auditory computer interface.

26 P. Fung, (November 1987) Novice Prolog Programmers.

27 P. Fung, B. DuBoulay & M. Elsom-Cook, (November 1987), An initial taxonomy of novices' misconceptions of the Prolog interpreter.

28 G. Kirkup, (November 1987), Considering the effect on women students of an increased use of microcomputers in distance education.

29 Sara Hennessy, Rick Evertsz, Dave Ellis, Phil Black, Tim O'Shea, Ann Floyd, Design Specification for 'Shopping on Mars' a computer-based Educational Activity.

30 A D N Edwards, The Use of home computers by disabled students at the Open University. Part 1: Previous use of computers in courses

31 Michael E. Fox (March 1988) A Report on Studies of Motivation, Teaching and Small Group Interaction with special reference to Computers and to the Teaching and Learning of Arithmetic.

32 Michael E. Fox (March 1988) Theory and Design for a Visual Calculator for Arithmetic.

33 M. Baker (February 1988) A Cognitive Model for Perception of Musical Grouping Structures

34 M. Baker (February 1988) An Artifical Intelligence Approach to Musical Grouping Analysis

35 P. Fung (December 1987) Novices' predictions of Prolog's control flow: A report on an empirical study

36 M. Elsom-Cook and F. Spensley (April 1988) Knowledge representation in a tutoring system for procedural skills.

37 F. Spensley & M. Elsom-Cook (April 1988) Dominie: Teaching and Assessment Strategies.

38 M. Baker (May 1988) Tutoring with Incomplete and Uncertain Knowledge